Asymptotic Capacity Bounds for Wireless Networks with Non-Uniform Traffic

Stavros Toumpis

Abstract

We develop bounds on the capacity of wireless networks when the traffic is non-uniform, i.e., not all nodes are required to receive and send similar volumes of traffic. Our results are asymptotic, i.e., they hold with probability going to unity as the number of nodes goes to infinity. We study *(i)* asymmetric networks, where the numbers of sources and destinations of traffic are unequal, *(ii)* multicast networks, in which each created packet has multiple destinations, *(iii)* cluster networks, that consist of clients and a limited number of cluster heads, and each client wants to communicate with any of the cluster heads, and *(iv)* hybrid networks, in which the nodes are supported by a limited infrastructure. Our findings quantify the fundamental capabilities of these wireless networks to handle traffic bottlenecks, and point to correct design principles that achieve the capacity without resorting to overly complicated protocols.

Keywords: Asymmetric traffic, capacity, clustering, hybrid networks, infrastructure support, mobile ad hoc networks, multicast routing, wireless access.

I. Introduction

We study the setting in which nodes equipped with wireless transceivers communicate over a shared wireless channel to create a multihop network. In this context, we develop bounds on the capacity of the network, which is defined as the theoretical limit on the total traffic that the network can support, assuming optimal coordination among the nodes. The bounds are determined assuming a number of different non-uniform traffic models under which some nodes are required to either create or receive much more traffic than other nodes. Following the approach introduced in [3], our results are asymptotic, i.e., they hold with probability going to unity as the number of nodes goes to infinity.

The author is with the Department of Electrical and Computer Engineering of the University of Cyprus, Kallipoleos 75, P.O. Box 20537, 1678 Nicosia, Cyprus. Part of this work was conducted while the author was at the Telecommunications Center Vienna (ftw.), in Vienna, Austria. Work was supported by **K***plus* funding for the ftw. project I0 “Signal and Information Processing”. A preliminary version of Sections IV, VI, and VII appeared in [1]. The basic result of Section IV also appeared, with sketches of proofs, in [2]. The basic result of Section V appeared, with a sketch of the proof, in the oral presentation of [1].

In [3], the authors consider a set of n nodes randomly placed on the surface of a sphere. Each of the nodes chooses another node as the destination for its traffic, randomly, uniformly and independently, and all n traffic streams are assumed to have a common rate requirement. The authors aim to find the maximum possible rate per stream $\lambda(n)$ that the network can achieve. Note that, because the placement of the nodes and the choice of destinations are random, $\lambda(n)$ is a random variable. The authors show that **with high probability (w.h.p.)**, i.e. with probability going to 1 as the number of nodes n goes to infinity, $\frac{K_1}{\sqrt{n \log n}} < \lambda(n) < \frac{K_2}{\sqrt{n \log n}}$, for some $K_2 > K_1 > 0$. Therefore, the maximum possible aggregate throughput $n\lambda(n)$ is **on the order of**, i.e., ignoring poly-logarithmic factors of the form $k_1(\log n)^{k_2}$, the square root of the nodes $\sqrt{n}$. Many researchers have followed the same tangent, and a significant number of results of the same flavor have accumulated [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14].

The traffic pattern used in [3] and almost all of the following works is, in some sense, as simple as possible: All nodes create data with the same rate $\lambda(n)$, and each of them picks at random one of the rest of the nodes as the destination for these data. For lack of a better description, we call this traffic pattern **uniform**. The uniform traffic pattern is a good model for certain networks, for example those used to support unicast voice transmission. On the other hand, there is a host of applications in which the traffic patterns will be fundamentally different. For example, in a network designed to support multimedia traffic between soldiers in a battlefield most of the traffic will have multiple destinations. As another example, in typical wireless sensor networks, a large number of sensors is interested in communicating with a relatively small number of sinks. With a few exceptions [6], [7], [8], [10], [11], the asymptotic properties of the capacity under such non-uniform traffic patterns remain to a large extend unexplored.

A. Contributions

In this work, we study wireless networks with no less than four different non-uniform traffic patterns which collectively cover a wide variety of scenarios. In particular, we calculate bounds on the capacity, i.e., the maximum possible aggregate throughput that the network can support under an optimal coordination of the nodes. Following the approach of [3], our results are asymptotic, i.e., they are only true with probability going to 1 as the number of nodes n goes to infinity. However, in contrast to [3], we achieve our results using basic tools of probability and a simple methodology, introduced in [14], which determines the rate with which the probability converges to unity. We also use a realistic channel model, that includes a general model for flat fading. In the interest of brevity, we focus on constructive lower bounds on the

capacity of networks, and formally derive only a few upper bounds. In addition, we also present a few other upper bounds with no formal proofs, but with strong heuristic justifications.

We first study **asymmetric networks**. These consist of two types of nodes: n source nodes, and n^d destination nodes[1], where $0 < d < 1$. Sources create packets with a common data rate, and the packets must be delivered to a single one of the destinations, chosen at random. Our main find is that when $d < \frac{1}{2}$, there are so few destinations, that bottlenecks start to form around them, constraining the maximum possible aggregate throughput to be around n^d. If, on the other hand, $d > \frac{1}{2}$, bottlenecks can be avoided, and the capacity is on the order of $n^{\frac{1}{2}}$, as in the uniform traffic setting of [3].

We then consider **multicast networks**. These consist of n nodes, each creating packets with a common data rate. Each packet must be delivered to n^d distinct nodes, chosen randomly among the rest. In this context, the capacity is on the order of $n^{\frac{d+1}{2}}$, and can be achieved without any multicasting in the media access layer, and using a multicast routing tree that can be constructed using only local information.

We also study **cluster networks**, which consist of n cluster nodes, and n^d cluster heads, where $0 < d < 1$. Each cluster node is the source of a traffic stream and the sink of a traffic stream. The traffic must be exchanged with *any* of the cluster heads, and all traffic streams have a common data rate. We show that the maximum possible aggregate throughput is on the order of n^d, and can be achieved (up to the order) without routing along multiple hops.

We conclude by studying **hybrid networks**, continuing the work of [6], [7]. These consist of n wireless nodes and n^d access points, where $0 < d < 1$. Access points are equipped with wireless transceivers and, in addition, are connected with each other through an independent network of infinite capacity. Each of the wireless nodes is creating traffic destined for one of the other wireless nodes, chosen at random. The access points have no traffic requirements of their own, but are there to support the communication of the wireless nodes. In this setting, we find that if $d < \frac{1}{2}$, then there are so few access points, that if the wireless nodes attempt to use them bottlenecks will be created. In that case, it is best for the wireless nodes to ignore the presence of the access points, and communicate with each other exclusively over the wireless channel. Therefore, the capacity is on the order of $n^{\frac{1}{2}}$. If, however, $d > \frac{1}{2}$, there is a sufficient number of access points to make a difference, and the capacity is on the order of n^d.

[1]Note that formally n^d must be an integer, which only occurs for certain combinations of n and d. However in the following we will ignore this and similar issues, as a more formal treatment, for example by using $\lfloor n^d \rfloor$, i.e., the integer part of n^d, would encumber the notation without in the least affecting the essence of the derivations.

The rest of this paper is organized as follows: in Section II we specify our network models and formally present our results. In Section III we present three lemmas that will be used throughout the text. Proofs for the results for asymmetric, multicast, cluster and hybrid networks are developed in Sections IV, V, VI, and VII respectively. We conclude in Section VIII.

II. Network Models and Results

Nodes are equipped with transceivers used for communication over a wireless channel of bandwidth W, and cannot transmit and receive simultaneously. Each node Z_i can transmit with any power $P_i \leq P_0$, where P_0 is a global maximum. When Z_i transmits with power P_i, Z_j receives the transmitted signal with power $G_{ij}P_i$, where $G_{ij} = Kf_{ij}|Z_i - Z_j|^{-\alpha}$. K is a constant, the same for all nodes, $|Z_i - Z_j|$ is the distance between nodes Z_i and Z_j, $\alpha > 2$ is the **decay exponent**, and the factor f_{ij} is the **fading coefficient**, a non-negative random variable that models fading.

We assume that the expectation $E[f_{ij}] = 1$, and that $f_{ij} = f_{ji}$. Distinct fading coefficients are independent and identically distributed (iid). We also assume that:

$$F^c(x) \triangleq P[f_{ij} > x] \leq \exp[-qx] \; \forall x > x_1, \tag{1}$$

for some $q,\ x_1 > 0$. Also, we assume that there is a median value $f_m > 0$ such that $P[f_{ij} \geq f_m] \geq \frac{1}{2}$. Both of these assumptions are satisfied by most distributions used to model fading, for example the Nakagami, Ricean and Rayleigh distributions, and the trivial distribution for which $P[f_{ij} = 1] = 1$.

Let $\{Z_t : t \in \mathcal{T}\}$ be the transmitting nodes at a given time, node Z_t transmitting with power P_t. Let us assume that node Z_j, $j \notin \mathcal{T}$ is receiving a data packet from Z_i, $i \in \mathcal{T}$. Then the **signal to interference and noise ratio (SINR)** at node Z_j will be $\gamma_j = \frac{G_{ij}P_i}{\eta + \sum_{k\in\mathcal{T},\ k\neq i} G_{kj}P_k}$, where η is the receiver thermal noise power, same for all nodes. The transmission will be successful if and only if, for the whole period of transmission, the transmission rate used, R_j, satisfies the inequality $R_j \leq f_R(\gamma_j) \triangleq W\log_2(1 + \frac{1}{\Gamma}\gamma_j)$. With $\Gamma > 1$, the equation approximates the maximum rate that meets a given BER requirement under a specific modulation and coding scheme [15]. With $\Gamma = 1$, it gives the Shannon bound.

A. Asymmetric Networks

Asymmetric networks consist of n **source nodes** X_1, X_2,. . ., X_n, and $m(n) = n^d$ **destination nodes** Y_1, Y_2, . . ., Y_m, placed randomly, uniformly and independently, in the unit square $\{(x, y) : |x|, |y| \leq \frac{1}{2}\}$.

We call $d \in (0,1)$ the **destination exponent**. Each source node is creating data traffic with a fixed data rate $\lambda(n)$ bps, common for all sources, that must be delivered to one of the destination nodes. Each source selects its destination randomly, uniformly, and independently of the others. Both types of nodes are allowed to transmit and receive, as well as relay packets.

The fundamental difference of this network from previously considered networks, such as the one in [3], is not that there are two types of nodes (sources and destinations), but the fact that their numbers n and $m(n)$ are different, and so the traffic is asymmetric: on the average more packets must arrive at each destination, than there are leaving each source. In fact, we could have just as well assumed that there are n destination nodes and only $m(n) = n^d$ source nodes, and arrived at very similar results. Applications where traffic asymmetries are expected are, for example, vehicular ad hoc networks in which many users will be downloading infotainment from a few central locations, and wireless sensor networks where the sensor nodes will be exchanging data with a small number of sinks.

We define the **capacity** $C(n)$ **of the network** as the supremum of all rates $\lambda(n)$ that are uniformly achievable by all sources, multiplied by their number n. Since the locations of the nodes, the destination of each data stream, and the fading coefficients are random, the capacity is a random variable.

Theorem 1: In asymmetric networks the capacity $C(n)$ is bounded w. h. p. as follows:

$$\left[\frac{4\alpha W}{\log 2}\right] n^d \log n \geq C(n) \geq \left[\frac{3\alpha-6}{3\alpha-5}\right]\left[\frac{Wqf_m 5^{-\frac{\alpha}{2}}}{676\Gamma\log 2}\right] \times \begin{cases} \frac{2}{27}\frac{n^{\frac{1}{2}}}{(\log n)^{\frac{3}{2}}} & \textit{if } \frac{1}{2} < d < 1, \\ \left[\frac{1-2d}{5}\right]\frac{n^d}{\log n} & \textit{if } 0 < d < \frac{1}{2}. \end{cases} \tag{2}$$

When $d < \frac{1}{2}$, bottlenecks are formed around the destinations, limiting the capacity of the network. If, however, $\frac{1}{2} < d < 1$, no bottlenecks are formed around the destinations, and the capacity can increase as fast as $n^{\frac{1}{2}}$, as in the uniform traffic case of [3]. Although we do not formally prove the upper bound $n^{\frac{1}{2}}$ on the capacity for the case $d > \frac{1}{2}$, it is intuitively clear from the work in [3] that it holds, and so the lower bound is always tight up to a poly-logarithmic factor of the form $k_1(\log n)^{k_2}$.

An important practical implication of Theorem 1 is that networks can handle well *some* asymmetry in the traffic, but designers should avoid any *extreme* asymmetry. In particular, the number of destinations $m(n)$ should be at least on the order of $n^{\frac{1}{2}}$, where n is the number of sources. For applications in which $m(n)$ is a design parameter, and it is useful to minimize it (because, for example, destinations are more expensive) the network has a ‘sweet-spot’: $m(n)$ should be around $n^{\frac{1}{2}}$. Using more destinations will not

improve the performance significantly, but using fewer will severely reduce it.

B. Multicast Networks

Multicast networks consist of n wireless nodes X_1, X_2,. . ., X_n, placed randomly, uniformly and independently, in the area $\{(x,y) : |x|,|y| \leq \frac{1}{2}\}$. Each node creates traffic with a common rate $\lambda(n)$ that is intended for $m(n) = n^d$ other nodes, that are chosen randomly, uniformly and independently among the rest. We call $d \in (0,1)$ the **multicast exponent**. Examples of networks with multicast traffic are wireless networks for military or search-and-rescue operations where each user might want to communicate with an arbitrary subset of the other users.

We define the **capacity** $C(n)$ **of the network** as the supremum of all rates $\lambda(n)$ that are uniformly achievable by all sources in the network, multiplied by their number n *and* the number of destinations $m(n) = n^d$. Note that the capacity is again a random variable.

Theorem 2: In multicast networks the capacity is bounded w. h. p. as follows:

$$C(n) \geq \left[\frac{3\alpha - 6}{3\alpha - 5}\right] \left[\frac{Wqf_m 5^{-\frac{\alpha}{2}}}{2700\Gamma \log 2}\right] \frac{n^{\frac{d+1}{2}}}{(\log n)^{1+\frac{d}{2}}}. \tag{3}$$

The improvement on the capacity over the uniform case is due to the possibility for the routing of each packet along a tree that passes through all destinations, as opposed to sending the same packet individually to each destination, in an uncoordinated manner. Although we formally present only a lower bound, we will use intuitive arguments to show that the routing tree employed by the constructive lower bound is of the same order of length as the minimum length multicast tree that the source can employ. For this reason, the lower bound is tight up to a poly-logarithmic factor.

It is reasonable that routing along efficient multicast trees can push the capacity of a network past the $n^{\frac{1}{2}}$ bound. Theorem 2 quantifies this fact. An interesting side result is that the tight lower bound can be achieved without employing multicasting on the media access layer. The intuitive justification of this rather unexpected result is that any efficient multicast trees will have such a small number of bifurcations, so that employing multicasting in the media access layer can not change the order of the capacity.

Another interesting side result is that the tight lower bound can be achieved without the source discovering the location of the destinations, or the destinations discovering the location of the source.

The only requirement is that each destination be discovered by a node carrying the packets that is on a distance at most $n^{-\frac{d}{2}}$ away from that destination.

C. Cluster Networks

Cluster networks consist of n **client nodes** X_1, X_2,. . ., X_n, and $m(n) = n^d$ **cluster heads** Y_1, Y_2, . . ., Y_m, placed randomly, uniformly and independently, in the area $\{(x, y) : |x|, |y| \leq \frac{1}{2}\}$. We call $d \in (0, 1)$ the **cluster head exponent**. Each client wants to establish a bidirectional communication (with rate $\lambda(n)$ in each direction) with *any* of the cluster heads. This model approximates well the traffic patterns that exist in wireless networks that operate using hierarchical clustering protocols, as for example Bluetooth [16]. Another application are sensor networks that consist of sensors and fusion centers.

We define the **capacity** $C(n)$ **of the network** as the supremum of all rates $\lambda(n)$ that are uniformly achievable by all data streams in the network, multiplied by their number $2n$. As in the previous cases, the capacity is a random variable.

Theorem 3: In cluster networks the capacity is bounded w. h. p. as follows:

$$\left[\frac{4\alpha W}{\log 2}\right] n^d \log n \geq C(n) \geq \left[\frac{Wqf_m}{338 \log 2\Gamma}\right] \left[\frac{3\alpha - 6}{3\alpha - 5}\right] 5^{-\frac{\alpha}{2}} \frac{n^d}{(\log n)^2}. \tag{4}$$

The theorem shows that, ignoring poly-logarithmic factors, the capacity increases with n roughly as n^d. The upper bound of (4) comes from the need of the network to share the area around the cluster heads. Therefore, the larger d is, the faster the capacity increases with n. In the context of networks that use clustering, the theorem suggests that, to maximize capacity, the size of clusters must be bounded, and so their number should increase linearly with n. If network designers are not willing to accept such a large number of clusters, they should be ready to sacrifice part of the capacity. The exact tradeoff is very simple, and is captured by Theorem 3. In the context of networks where the cluster heads are gateways to the outside world, the theorem suggests that there is no limit to how many gateways are needed: the greater the investment of the network provider (i.e., the larger d is), the larger the capacity is going to be. Again, the tradeoff is very simple and is captured by Theorem 3.

Finally, as the proof will show, the lower bound on the capacity can be achieved even if clients do not transmit to each other, but provided they are not restricted to communicate with the nearest cluster head. In other words, advanced routing protocols can not change the capacity by more than a poly-logarithmic

factor, and designers should focus instead on efficient polling algorithms and the efficient handling of bottlenecks around the cluster heads.

D. Hybrid Networks

Hybrid networks consist of n **wireless nodes** X_1, X_2,. . ., X_n, and $m(n) = n^d$ **access points** Y_1, Y_2, . . ., Y_m, placed randomly, uniformly and independently, in the two-dimensional area $\{(x, y) : |x|, |y| \leq \frac{1}{2}\}$. We call $d \in (0, 1)$ the **access point exponent**. We assume that the access points are connected with each other through a data link of infinite capacity that does not consume any of the available bandwidth W. There are n traffic streams and each wireless node is the source of a single stream, and the destination of a single stream. A node cannot be the source and destination of the *same* stream. Apart from this restriction, all other combinations of sources and destinations are equally probable. The access points do not have any communication needs of their own, but are there to support the wireless nodes.

This network shares important common characteristics with both pure wireless multihop networks and also pure cellular networks: On the one hand, it partly consists of a large number of wireless nodes that communicate over a wireless channel and can route each other's traffic, as in wireless multihop networks. On the other hand, the wireless nodes are supported by access points that form an independent network with infinite capacity and do not have any traffic needs of their own; their role is similar to that of base stations in cellular networks. The asymptotic capacity of such networks was first studied in [6], [7], and is of great practical interest, as future generation cellular systems will be using this hybrid topology.

We define the **capacity** $C(n)$ **of the network** as the supremum of all rates $\lambda(n)$ that are uniformly achievable by all data streams in the network, multiplied by their number n.

Theorem 4: In hybrid networks the capacity is bounded w. h. p. as follows:

$$C(n) \geq \frac{1}{2}\left[\frac{Wqf_m}{338\log 2\Gamma}\right]\left[\frac{3\alpha-6}{3\alpha-5}\right]5^{-\frac{\alpha}{2}}\frac{n^d}{(\log n)^2}, \tag{5}$$

$$C(n) \geq \left[\frac{10^{-\frac{\alpha+3}{2}}}{648}\frac{3\alpha-6}{3\alpha-5}\frac{Wqf_m}{\Gamma}\right]n^{\frac{1}{2}}(\log n)^{-\frac{3}{2}}. \tag{6}$$

Although we do not formally prove upper bounds, we provide an intuitive justification that (5) is tight when $d > \frac{1}{2}$, and (6) is tight when $d < \frac{1}{2}$.

The theorem suggests that more than $n^{\frac{1}{2}}$ access points are needed for the infinite-capacity infrastructure to have any effect on the performance of the network. As the proof will reveal, when $d < \frac{1}{2}$, the wireless

nodes should ignore the available infrastructure, and use each other for communication. That way, an aggregate throughput on the order of $n^{\frac{1}{2}}$ is achievable. If, however, $\frac{1}{2} < d < 1$, the wireless nodes should not depend on each other for routing their traffic, but rather should make heavy use of the infrastructure.

We note that a similar result was first reported in [6], [7]. Our setup, however, is different in a number of critical ways: Firstly, we require that all wireless nodes are guaranteed the same throughput. Secondly, the locations on the access points are random, and finally we assume a more realistic channel model, that includes a general fading model.

III. Useful Lemmas

The first lemma is closely related to the well-known Coupon Collector's Problem [17], however, to the best of our understanding, it has not appeared elsewhere in this form.

Lemma 1: Let n balls be placed in m urns, uniformly and independently of each other. Let b_j, $j = 1, \ldots, m$ be the number of balls that end up in the j-th urn. Then for any $\epsilon > 0$ there is a $\delta(\epsilon) > 0$ such that $P[\forall j \ (1-\epsilon)\frac{n}{m} \leq b_j \leq (1+\epsilon)\frac{n}{m}] \geq 1 - 2m\exp[-\delta(\epsilon)\frac{n}{m}]$.

Proof: We make use of Chernoff's bounds [18]: Let X be a binomial random variable, with parameters k (the number of experiments) and p (the probability of success of each experiment). For any $\epsilon > 0$,

$$P[X < (1-\epsilon)kp] \quad < \quad \exp[-kp\frac{\epsilon^2}{2}], \tag{7}$$

$$P[X > (1+\epsilon)kp] \quad < \quad \frac{\exp[\epsilon kp]}{(1+\epsilon)^{(1+\epsilon)kp}} \triangleq \exp[-kpf(\epsilon)], \tag{8}$$

where $f(\epsilon) \triangleq (1+\epsilon)\log(1+\epsilon) - \epsilon$. By calculating its derivative, we have that $f(\epsilon) > 0$ for $\epsilon > 0$.

Since each ball is placed in an urn independently of the others, b_j follows the binomial distribution, with number of experiments equal to n and probability of success equal to $\frac{1}{m}$. (Note, however, that the b_j are not independent.) Applying Chernoff's bounds, we have:

$$P[b_j < (1-\epsilon)\frac{n}{m}] < \exp[-\frac{\epsilon^2}{2}\frac{n}{m}], \qquad P[b_j > (1+\epsilon)\frac{n}{m}] < \exp[-f(\epsilon)\frac{n}{m}]. \tag{9}$$

We note the basic inequality $P(\cup_{j=1}^m E_j) \leq \sum_{j=1}^m P(E_j)$, typically referred to as the *union bound*. Then:

$$\begin{aligned} P[\forall j \ (1-\epsilon)\frac{n}{m} \leq b_j \leq (1+\epsilon)\frac{n}{m}] &\geq 1 - \sum_{j=1}^{m} \left\{ P[b_j < (1-\epsilon)\frac{n}{m}] + P[b_j > (1+\epsilon)\frac{n}{m}] \right\} \\ &\geq 1 - m\left\{ \exp[-\frac{\epsilon^2}{2}\frac{n}{m}] + \exp[-f(\epsilon)\frac{n}{m}] \right\} \geq 1 - 2m\exp[-\delta(\epsilon)\frac{n}{m}], \end{aligned}$$

where $\delta(\epsilon) \triangleq \max\{\frac{\epsilon^2}{2}, f(\epsilon)\} > 0$. The first inequality comes from the union bound, and the second inequality from (9). □

In subsequent sections, we will have to bound the effects of interfering transmissions in the reception of signals. As the fading distribution has an exponentially thin tail, the following lemma applies:

Lemma 2: Let n *nodes. W. h. p.,* $\max_{1\le i<j\le n}\{f_{ij}\} \le \frac{3}{q}\log n$.

Proof: Let the events $F_{ij}(n) \triangleq \{f_{ij} > \frac{3}{q}\log n\}$. Then:

$$P\left[\max_{1\le i<j\le n}\{f_{ij}\} \le \frac{3}{q}\log n\right] = 1 - P[\cup_{1\le i<j\le n}F_{ij}(n)] \ge 1 - \sum_{1\le i<j\le n} P[F_{ij}(n)] \ge 1 - \frac{n(n-1)}{2}n^{-3} \to 1.$$

The first inequality comes from the union bound. The second comes from symmetry and applying (1), and holds only for sufficiently high n. □

Finally, observe that if a sequence of events A_n occurs w. h. p., and a second sequence of events B_n occurs w. h. p. conditioned on the sequence A_n, then B_n will also occur w. h. p. without the conditioning:

Lemma 3: Let $\lim_{n\to\infty} P(A_n) = 1$ *and* $\lim_{n\to\infty} P(B_n|A_n) = 1$. *Then* $\lim_{n\to\infty} P(B_n) = 1$.

The proof follows immediately by noting that $P(B_n) = P(B_n|A_n)P(A_n) + P(B_n|A_n^c)P(A_n^c)$. In practical terms, if we need to prove that a sequence of events occurs w. h. p., we are free to condition the discussion on any sequence of events that occurs also w. h. p. It is also clear that we can iteratively condition on more than one sequence of events.

IV. Asymmetric Networks

In this section we develop a constructive proof for Theorem 1 in the spirit of [3]: we develop a communications scheme whose aggregate throughput equals the lower bound w. h. p., and as the capacity is the supremum of the aggregate throughputs of *all* schemes, it will necessarily exceed this lower bound.

A. Cell Lattice

As shown in Fig. 1, we divide the square region $\{(x,y) : |x|,|y| \le \frac{1}{2}\}$ in a regular lattice of $g(n) = \frac{n}{9\log n} \triangleq r^2$ cells $c_1,\ c_2,\ \ldots, c_{g(n)}$. Each cell can be identified by its coordinates (v_1, v_2) in the lattice, where $1 \le v_1,\ v_2 \le r$; the cell on the lower left corner has coordinates $(1,1)$. We call two cells **neighbors** if they share a common boundary edge, so that each cell has at most four neighbors.

Let s_j be the number of source nodes in cell c_j. Thinking of cells like urns and source nodes like balls, we see that Lemma 1 applies. Setting $\epsilon = \frac{1}{2}$, $m = g(n)$, $b_j = s_j$, $\delta(\epsilon) = \max\{\frac{\epsilon^2}{2}, (1+\epsilon)\log(1+\epsilon)-\epsilon\} = \frac{1}{8}$, it follows that $P\left[\forall j \ \frac{9\log n}{2} \le s_j \le \frac{27\log n}{2}\right] \ge 1-\frac{2n^{-\frac{1}{8}}}{9\log n}$, which goes to 1 as $n \to \infty$. Therefore, by Lemma 3 we take the following to always hold:

$$\forall j \ \frac{9\log n}{2} \le s_j \le \frac{27\log n}{2}. \tag{10}$$

Next, let $F_{ij}(n)$ be the event that the source node X_i can not find a source node in one of its neighboring cells c_j, such that their mutual fading coefficient is greater or equal to f_m. By the independence of the fading coefficients, and using (10), it follows that $P[F_{ij}(n)] \le (\frac{1}{2})^{\frac{9\log n}{2}}$. By using the union bound, and noting that there are n source nodes, each with at most 4 neighboring cells, it follows that $P[\cup_{i,j}F_{ij}] \le 4n(\frac{1}{2})^{\frac{9\log n}{2}} \to 0$. Therefore, w. h. p. each source node will be able to find another source node in each of the neighboring cells, such that their mutual fading coefficient is equal to or greater than f_m.

Finally, let $G_{ij}(n)$ be the event that a destination node Y_i and a source node X_j lying in the same cell will not be able to find a relaying source node X_k, also on that cell, such that the mutual fading coefficients $f_{Y_iX_k} \ge f_m$ and $f_{X_kX_j} \ge f_m$. By the independence of the fading coefficients, the probability that a particular source node cannot be used is at most $\frac{3}{4}$, and the probability that there is no source node that can be used is at most $(\frac{3}{4})^{\frac{9}{2}\log n}$. Applying the union bound, it follows that the probability $P[\cup_{i,j}G_{ij}(n)] \le n^d(\frac{27}{2}\log n)(\frac{3}{4})^{\frac{9}{2}\log n} \to 0$. Therefore, w.h.p. any destination node will be able to communicate with any source node in its cell, by using another source node in that cell as a relay, and in both hops the fading coefficient will be greater or equal to the median f_m.

Let us summarize the results until now: We have divided our area into $\frac{n}{9\log n}$ cells and we have shown that the following properties hold w. h. p., and so can taken from granted, using Lemma 3 for justification: **(i)** The numbers of source nodes in all cells are bounded by (10). **(ii)** Each source node can find a source node in any of its neighboring cells so that their mutual fading coefficient is greater than or equal to the median f_m. **(iii)** Each source node can communicate with any of the destination nodes in its cell through a relaying source node in that cell, so that the fading coefficients of both hops are greater than or equal to the median f_m.

B. Routing Protocol

As shown in Fig. 2, packets are routed according to the following rules:

(i) If a source node X_j has data packets (possibly not created at X_j) that must be delivered to a destination node Y_i lying in the same cell, and $f_{X_jY_i} < f_m$, X_j will transmit the data packets to another source node X_k lying in the same cell, for which $f_{X_jX_k} \geq f_m$ and $f_{X_kY_i} \geq f_m$. Node X_k will then transmit the packet to the destination node Y_i. By the discussion of Section IV-A, we can assume that such a node exists.

(ii) If the destination node Y_j of a source node X_i lies in a different cell from X_i, the packets of X_i are routed through intermediate cells. In particular, only communication between source nodes who lie in neighboring cells and whose mutual fading coefficient is at least equal to the median is allowed. In addition, the packets are first transmitted along cells whose x-coordinate is the same as the x-coordinate of the source, until they arrive at a cell whose y-coordinate is the same as the y-coordinate of the destination. Then, the packets are transmitted along cells whose y-coordinate is the same as the y-coordinate of the destination, until they arrive at a source node lying in the same cell with the destination. By the discussion of Section IV-A, we can assume that such relays always exist. Once the packets arrive at the cell of the destination, they are delivered to the destination as specified by rule **(i)**.

To evaluate the performance of this scheme, we must calculate the load that the routing protocol creates for each cell. To this end, let us define M_j as the number of source nodes that lie in cells whose x-coordinate is the same as the x-coordinate of cell c_j, and N_j as the number of destination nodes that lie in cells whose y-coordinate is the same as the y-coordinate of cell c_j. We develop bounds on the values of M_j and N_j that we will use to bound the traffic that each cell must support.

To bound the value of M_j, we note that there are $\sqrt{\frac{n}{9\log n}}$ cells with the same x-coordinate with cell c_j, each with at most $\frac{27}{2}\log n$ source nodes. Therefore:

$$\forall j,\ \ M_j \leq \frac{9}{2}\sqrt{n\log n}. \tag{11}$$

Next, we bound N_j for the case $d > \frac{1}{2}$. Applying Lemma 1 where the balls are the destination nodes and the urns are the rows:

$$P[\forall j\ \ \frac{3}{2}n^{d-\frac{1}{2}}\sqrt{\log n} \leq N_j \leq \frac{9}{2}n^{d-\frac{1}{2}}\sqrt{\log n}] \geq 1 - 2\sqrt{\frac{n}{9\log n}}\exp[-\delta(\frac{1}{2})3n^{d-\frac{1}{2}}\sqrt{\log n}] \to 1.$$

By Lemma 3, we are allowed to assume that:

$$d > \frac{1}{2} \Rightarrow \forall j,\ \ N_j \leq \frac{9}{2}n^{d-\frac{1}{2}}\sqrt{\log n}. \tag{12}$$

Finally, we uniformly bound the N_j for the case $d < \frac{1}{2}$. For this we use (8), noting that N_j follows the binomial distribution with $p = (\frac{n}{9\log n})^{-\frac{1}{2}}$ and $k = n^d$. Setting ϵ to satisfy $(1+\epsilon)kp = x$, where x

will be specified later, we have that: $P[N_j > x] < \exp[x - kp](\frac{kp}{x})^x < \frac{\exp[x]}{x^x}\left(3n^{d-\frac{1}{2}}\sqrt{\log n}\right)^x$. Applying the union bound, we have that $P[\exists j: \ N_j > x] \leq \frac{\exp[x]}{x^x}\left(3n^{d-\frac{1}{2}}\sqrt{\log n}\right)^x \sqrt{\frac{n}{3\log n}}$, which goes to 0 if we choose $x > \frac{1}{1-2d}$, for example $x = \frac{2}{1-2d}$. Applying Lemma 3, we can assume that:

$$d < \frac{1}{2} \Rightarrow \forall j, \ N_j \leq \frac{2}{1-2d}. \tag{13}$$

Lemma 4: Let r_j be the number of routes arriving, and possibly terminating, at cell c_j. Then w. h. p.:

$$\forall \ j, \ r_j \leq r_{\max}(n) \triangleq \begin{cases} \frac{27}{2}(n\log n)^{\frac{1}{2}} & \textit{if} \ \ \frac{1}{2} < d < 1, \\ \frac{5}{1-2d}n^{1-d} & \textit{if} \ \ 0 < d < \frac{1}{2}. \end{cases}$$

Proof: Let r_{j1} be the number of routes that cross c_j while on their vertical leg (see Fig. 2). The sources of those routes share a common x-coordinate with c_j. Also, let r_{j2} be the number of routes that cross c_j while on their horizontal leg. The destination nodes of these routes share a common y-coordinate with c_j. Each route crossing c_j will belong to one or both of the two types of routes, so necessarily $r_j \leq r_{j1} + r_{j2}$. Therefore, it suffices to bound both r_{j1} and r_{j2} uniformly for all cells c_j.

As each node is the source of a single stream, $r_{j1} \leq M_j$. To bound r_{j2}, we note that, by a straightforward application of Lemma 1, at most $2n^{1-d}$ routes can be terminating at each destination, w. h. p. Therefore $r_{j2} \leq 2n^{1-d}N_j$ w. h. p. Combining these inequalities we have that $r_j \leq M_j + 2n^{1-d}N_j$ w. h. p., for all cells c_j. The result follows by using (11), (12), and (13), also noting that when $d < \frac{1}{2}$, $\frac{\sqrt{n\log n}}{n^{1-d}} \to 0$. $\square$

Since there are n routes, each requiring a number of hops on the order of $(\frac{n}{\log n})^{\frac{1}{2}}$, and the total number of hops must be shared by $\frac{n}{9\log n}$ cells, on the average each cell will be required to relay a number of routes on the order of $(n\log n)^{\frac{1}{2}}$. Therefore, Lemma 4 implies that when $d > \frac{1}{2}$, no cell will have to carry much more that its 'fair share' of the traffic. If, however, $d < \frac{1}{2}$, then there are so few destinations, that a few 'unlucky' cells (those on the same column with a destination) will be required to serve around n^{1-d} routes, which is much more than their 'fair share' of traffic. In those cells, bottlenecks will form.

C. Time Division

We divide the $g(n) = r^2$ cells into nine regular sub-lattices, such that any two cells belonging in the same sub-lattice are separated by at least two cells belonging to different sub-lattices. In Fig. 3 we have shaded the cells belonging to one of the 9 sub-lattices.

We divide time into frames, and each frame into nine slots, each slot corresponding to a sub-lattice. At any time during that slot, only one node from each cell of the corresponding sub-lattice is allowed to *receive* (but many nodes in that cell may receive consecutively in the same slot). Because of the way we constructed the routing protocol, the transmitter of that transmission will have to lie in the same cell, or in one of the four neighboring cells. All transmissions are with the maximum power P_0.

Lemma 5: The SINR γ_j at any source or destination node Z_j that is receiving is bounded w. h. p. by

$$\gamma_j > \gamma_{\min}(n) \triangleq 5^{-\frac{\alpha}{2}} \left[\frac{3\alpha-6}{3\alpha-5}\right] \left[\frac{qf_m}{25}\right] \frac{1}{\log n}. \tag{14}$$

Proof: We first bound the interference I_j. For this, we first note that by Lemma 2, w. h. p. no fading coefficient is greater than $\frac{3}{q}\log n$. Next, let $x_0 = \frac{1}{r}$ be the length of the sides of the cells, and let c_k be the cell in which the receiving node lies. Working as in [19], we note that the rest of the cells in the same sub-lattice are located along the perimeters of concentric squares, whose center is cell c_k. Irrespective of the coordinates of c_k, all the cells of its sub-lattice are located along the perimeters of at most $\lfloor \frac{r-1}{3} \rfloor$ squares. There are at most $8i$ interferers corresponding to the i-th square, whose distances from the receiver will be at least $x_0(3i-2)$. Consequently, the interference at the receiver is upper bounded by

$$\begin{aligned} I_j &\leq \left[\frac{3}{q}\log n\right] \sum_{i=1}^{\lfloor \frac{r-1}{3} \rfloor} \frac{8iKP_0}{[x_0(3i-2)]^\alpha} \leq \left[\frac{3}{q}\log n\right] \frac{8KP_0}{x_0^\alpha} \left[1 + \sum_{i=2}^{r} (3i-2)^{1-\alpha}\right] \\ &< \left[\frac{3}{q}\log n\right] \frac{8KP_0}{x_0^\alpha} \left[1 + \int_0^r (3x+1)^{1-\alpha}\, dx\right] \leq \left[\frac{3}{q}\log n\right] \frac{8KP_0}{x_0^\alpha} \left[\frac{3\alpha-5}{3\alpha-6}\right]. \end{aligned} \tag{15}$$

We also need a lower bound on the power of the useful signal. Clearly, since the maximum possible distance that the useful signal will need to travel, under the routing assumptions, is $\sqrt{5}x_0$, and the fading coefficient between the transmitter and the receiver is at least equal to f_m, w.h.p. we have that $S_j \geq KP_0 f_m(\sqrt{5}x_0)^{-\alpha}$. Combining this with (15), and noting that the thermal noise remains bounded, and therefore becomes negligible as $n \to \infty$, we arrive at (14). □

We now assume that all transmitters transmit with rate $f_R(\gamma_{\min})$. By Lemma 5, w. h. p. all transmissions will be successful.

D. Lower Bound

The nodes of each cell are allowed to receive during only 1 out of 9 slots, and with rate equal to $f_R(\gamma_{\min}(n))$. The number of routes that will be crossing each cell c_j is upper bounded by $r_{\max}(n)$,

determined by Lemma 4. Most of these routes will require one reception, however a few of these, in particular those whose destination lies in cell c_j, may require three receptions. Therefore, each route, and its associated source node, is guaranteed a rate of communication $\lambda(n) = f_R(\gamma_{\min})\left[3 \times 9 \times r_{\max}(n)\right]^{-1}$. Multiplying by n, and substituting for $r_{\max}(n)$ and $\gamma_{\min}(n)$ from Lemmas 4 and 5 respectively, we see that our scheme achieves an aggregate throughput equal to the lower bound of (2). Since the capacity is the supremum of the aggregate throughputs of *all* possible schemes, it will necessarily be greater than the aggregate throughput of our scheme, and the result follows.

E. Proof of Upper Bound

Let $d_{\min}$ be the minimum of all distances between all mn source-destination pairs, and let $H_{ij}(x)$ be the event $\{|X_i - Y_j| \le x\}$. Then:

$$P[d_{\min} \le x] = P[\cup_{i,j} H_{ij}(x)] \le \sum_{i=1}^{n}\sum_{j=1}^{m} P[H_{ij}(x)] = nmP[H_{11}(x)] \le nm\pi x^2. \tag{16}$$

The first inequality comes from the union bound. The second equality comes from using symmetry. The last inequality comes from noting that the nodes are placed in a square with surface area equal to 1, and that nodes X_1 and Y_1 will be within distance x of each other if Y_1 is placed on the intersection of the square with a disk of radius x, centered at node X_1.

The capacity is less than the aggregate throughput $T(n)$ that would have been achieved if all destination nodes were receiving, for all time, using the whole bandwidth, and without experiencing interference from competing transmissions. Therefore, we can bound $T(n)$ as follows:

$$T(n) \le mW\log_2\left(1 + \frac{1}{\Gamma}\frac{KP_0 d_{\min}^{-\alpha}\frac{3}{q}\log n}{\eta}\right) \le mW\log_2\left(1 + \frac{3KP_0}{\eta q\Gamma} n^{3\alpha}\log n\right) \le \left[\frac{4\alpha W}{\log 2}\right] n^d \log n.$$

The first inequality comes from assuming that all destinations receive all the time and with no interference, and by using the bound on the value of the fading coefficients of Lemma 2. The second inequality holds w.h.p., and comes by applying (16) with $x = n^{-3}$. The last holds for sufficiently large values of n, and comes using simple properties of the logarithm function. Since $C(n) \le T(n)$, the bound follows.

V. Multicast Networks

In order to better motivate the proof of the lower bound of Theorem 2, we first present a heuristic upper bound[2]: to minimize the number of transmissions needed for a packet to reach all its destinations, it is

[2]We note that a similar bound, based on an alternative heuristic argument, appeared in the independent work in [11].

clear that the packet must be routed along a **multicast tree** which passes through all the destinations and has as small a length as possible. Let us find a lower bound on the length of *any* tree that connects all destinations. For this, let us divide the whole region in $g(n) = \frac{n^d}{9 \log n}$ regular cells, each with side length equal to $g(n)^{-\frac{1}{2}}$. Working as Section IV, it follows that there will be a destination in each of them, so a tree connecting all of them will have a length on the order of $g(n) \times g(n)^{-\frac{1}{2}} = g(n)^{\frac{1}{2}} \simeq n^{\frac{d}{2}}$, ignoring logarithmic factors. Assuming transmissions across distances which are as small as possible, i.e., on the order of $n^{-\frac{1}{2}}$, it follows that each packet will need roughly $n^{\frac{d+1}{2}}$ transmissions to be delivered to all n^d destinations. As the number of simultaneous transmissions (across distances on the order of $n^{-\frac{1}{2}}$) over the whole network is on the order of n (using, for example the time division scheme of the previous section), it follows that the maximum possible aggregate rate of packet deliveries at destinations is on the order of $n \times n^d \times [n^{\frac{d+1}{2}}]^{-1} = n^{\frac{d+1}{2}}$. Up to the exponent of n, this upper bound equals the lower bound.

Moving to the constructive proof of the lower bound, ideally, we would like to construct a scheme that uses a multicast tree that is as short as possible, for example a Steiner tree on a properly defined graph. However, we also need a tree that is amenable to analysis. The tree we now specify represents a good compromise between these goals.

First, we divide the square region $\{(x, y) : |x|, |y| \leq \frac{1}{2}\}$ into $g(n) = \frac{n}{9 \log n}$ cells. The properties **(i)-(iii)** of Section IV-A continue to hold, where now property **(iii)** applies to the communication of two nodes in the same cell. As shown in Fig. 4, the tree consists of three legs:

First leg: The packet is propagated along a straight line to all cells which have the same y-coordinate as the cell of the source.

Second leg: Starting from the cell of the source, every $h(n) \triangleq \frac{n^{\frac{1-d}{2}}}{3\sqrt{\log n}}$ cells along the first leg, the packet also propagates along the vertical direction. Therefore, there are $n^{\frac{d}{2}}$ vertical legs per tree, separated by a distance of $n^{-\frac{d}{2}}$.

Third leg: Each destination receives the packet from the cell that received the packet in the second leg which is closest.

As with the routing protocol of Section IV-B, communication is between nodes those mutual fading coefficient is no smaller than the median f_m. Also, if a packet reaches a node in the cell of the destination other than the destination, it will reach the destination by two more hops, through a relay node, such that both mutual fading coefficients are at least equal to the median.

The aim of the first two legs is to spread the packet uniformly through the whole region, and the number of vertical sections strikes the optimal balance between having a small number of total hops and a thick coverage. Ignoring logarithmic factors and the fact that packets do not follow exactly straight lines, we note that the length of the tree is $1 + n^{\frac{d}{2}} \times 1 + n^{d} \times n^{-\frac{d}{2}} \simeq n^{\frac{d}{2}}$. Therefore, this tree has the potential to achieve our heuristic upper bound, at least up to a poly-logarithmic factor.

Next, we develop an upper bound on the traffic supported by each cell. For this, let r_j be the number of routes that cell c_j must support, and let r_{j1}, r_{j2}, and r_{j3} be the total number of routes passing through c_j in their first, second, and third leg respectively. Clearly, $r_j \leq r_{j1} + r_{j2} + r_{j3}$.

To bound r_{j1}, we apply Lemma 1 as in the case of the bound (11) and conclude that, for all j, $r_{j1} \leq \frac{9}{2}\sqrt{n \log n}$. To bound r_{j2}, we note that each of the n nodes will contribute to r_{j2} with one route with probability $h(n)^{-1}$, and so by a simple application of the Chernoff and union bounds, w.h.p., for all j, $r_{j2} \leq \frac{3}{2}h(n)^{-1}n = \frac{9}{2}n^{\frac{1+d}{2}}\sqrt{\log n}$. To bound r_{j3}, we note that a cell c_j will only have to serve *some* of the third legs of routes with destinations that lie in either c_j, or in one of the $h(n)$ cells on its left, or in one of the $h(n)$ cells on its right. Therefore, r_{j3} is at most equal to the number of destinations in $2h(n) + 1$ cells. By Lemma 1, w.h.p there are at most $\frac{27}{2}\log n$ nodes in each of these cells, for a total of at most $[2h(n)+1]\frac{27}{2}\log n$ nodes. The probability that one of these nodes is chosen when a node chooses his next destination is *at most* $[2h(n)+1]\frac{27}{2}(\log n)n^{-1}$ (it is smaller if one or more of the nodes among them have already been picked). Applying the Chernoff bound (8) with number of experiments n^{1+d} and probability of success $[2h(n)+1]\frac{27}{2}(\log n)n^{-1}$, it follows that $r_{j3} \leq \frac{27}{2}n^{\frac{1+d}{2}}\sqrt{\log n}$, with probability going to 1 exponentially fast. By a simple application of the union bound, it follows that w.h.p. the inequality will hold for all j. Combining the bounds for r_{j1}, r_{j2}, r_{j3}, it follows that w.h.p.,

$$\forall j,\ r_j \leq r_{\max}(n) \triangleq 4(n \log n)^{\frac{1+d}{2}}. \tag{17}$$

Next, we specify that the nodes use the time division schedule of Section IV-C, under which each receiver is guaranteed, w.h.p., an SINR equal to the bound $\gamma_{\min}(n)$ given by (14). Also, every transmitter transmits with rate $f_R(\gamma_{\min}(n))$. The number of routes crossing each cell is at most $r_{\max}(n)$, given by (17). Most of these involve just one hop, however those few whose destination lies in cell c_j will require three transmissions. Therefore, each route is guaranteed a rate of communication $\lambda(n) = f_R(\gamma_{\min}(n))\,[3 \times 9 \times r_{\max}(n)]^{-1}$. Multiplying with n, for the number of nodes, and n^d, for the number of destinations of each node, we arrive that the lower bound (3).

VI. Cluster Networks

Regarding the upper bound of Theorem 3, we simply note that we can prove it by applying the technique used for proving the upper bound of (2): we must simply consider upper bounds on the aggregate throughput received at the *cluster heads*, as opposed to the *destination nodes*.

We next present a constructive proof of the lower bound of (4). We divide the square region $\{(x,y):|x|,|y|\leq\frac{1}{2}\}$ in a regular lattice of $g(n)=\frac{n^d}{9\log n}\triangleq r^2$ cells, as shown in Fig. 1. Let s_j and d_j be the numbers of client nodes and cluster heads respectively in cell c_j. By Lemma 1, it follows that w. h. p.

$$\forall j,\quad \frac{9}{2}n^{1-d}\log n\leq s_j\leq\frac{27}{2}n^{1-d}\log n, \tag{18}$$

$$\forall j,\quad \frac{9}{2}\log n\leq d_j\leq\frac{27}{2}\log n. \tag{19}$$

The probability that a client will not be able to find a cluster head in its own cell such that their mutual fading coefficient is greater than the median f_m is, using the independence of different fading coefficients, at most $(\frac{1}{2})^{(\frac{9}{2}\log n)}$. Using the union bound, it follows that the probability that *any* of the clients will not be able to find such a cluster head is smaller than $n(\frac{1}{2})^{(\frac{9}{2}\log n)}$, which converges to 0 as $n\to\infty$. Therefore, w. h. p. all clients will have a fading coefficient to one of the cluster heads that is at least equal to f_m.

In addition, we impose on the nodes the time division scheme of Section IV-C: time is divided in frames, and each frame in 9 slots. At any time during a slot, only a single node (either a cluster head or a client node) from each cell of the corresponding sub-lattice is allowed to transmit, and with maximum power. Since the receiver necessarily lies in the same cell, the lower bound on the SINR of Lemma 5 continues to hold. Therefore, if the transmitter transmits with rate $f_R(\gamma_{\min}(n))$, where $\gamma_{\min}(n)$ is given by (14), w. h. p. all transmissions will be successful.

By (18), there are less than $\frac{27}{2}n^{1-d}\log n$ client nodes in each slot. We divide each slot in $2\times[\frac{27}{2}n^{1-d}\log n]$ time intervals, each of which is devoted to the transmission of a packet either from or to a client node. Each stream of data is guaranteed a rate of communication equal to $\lambda(n)=f_R(\gamma_{\min}(n))[27n^{1-d}\log n]^{-1}$. Multiplying by $2n$ for the total number of streams, and substituting for $\gamma_{\min}(n)$ from Lemma 5, we arrive at the lower bound of (4).

VII. Hybrid Networks

Because of the similarities between cluster and hybrid networks, the wireless nodes can use for their communication the scheme that was used in Section VI for proving the lower bound of (4). In particular,

wireless nodes do not transmit to each other, but rather transmit directly to an access point near by. The packet is then transmitted through the infinite capacity network to an access point close to its destination, and is then transmitted one more time through the use of the wireless interface to the destination. All the analysis of Section VI goes through, if we substitute client nodes with wireless nodes and cluster heads with access points. The only difference is that, because each packet must be transmitted twice, the aggregate throughput is half the throughput achieved in cluster networks. The bound (5) follows.

To derive (6), we consider the opposite extreme. In particular, we note that the n wireless nodes are free to ignore the infrastructure of the access points, and establish a communication scheme using only themselves. This case has been studied independently in [14], where it was shown (equation (4) of that work) that, under uniform traffic conditions, it is possible to achieve a per-node throughput equal to:

$$\lambda(n) = \left[\frac{10^{-\frac{a+3}{2}}}{648} \frac{3\alpha - 6}{3\alpha - 5} \frac{W q f_m}{\Gamma} \right] n^{-\frac{1}{2}} (\log n)^{-\frac{3}{2}}.$$

The derivation is omitted, as it appears in [14]. However, we note that it is very similar to (and in fact simpler than) the derivation leading to the lower bound of (2). Multiplying $\lambda(n)$ by the number of nodes n, we derive the lower bound of (6).

Although we provide no formal proof, it is intuitively clear that, in the case $d > \frac{1}{2}$, the bound (5) is tight, up to a poly-logarithmic factor. Indeed, the aggregate throughput of packets using the infrastructure, even for part of their transport, can not exceed the upper bound of (4), and the aggregate throughput of packets not using the infrastructure is much less, on the order of $n^{\frac{1}{2}}$, by [3]. By a similar argument, the bound (6) is tight, up to a poly-logarithmic factor, when $d < \frac{1}{2}$.

VIII. Conclusions

We study wireless networks with four different traffic models: asymmetric, multicast, cluster and hybrid. The common aspect of these models is their non-uniformity: in each of them some nodes are required to either send or collect much more traffic than other nodes This lack of uniformity places a strain on the network, through the formation of traffic bottlenecks that have the potential to reduce the capacity.

We present lower and upper bounds on the capacity that hold with probability going to unity as the number of nodes in the network goes to infinity. In the interest of brevity, we also present a number of conceptually straightforward upper bounds with only intuitive justification.

Our work quantifies the inherent capabilities of wireless networks to handle various types of traffic non-uniformities, and provides useful guidelines to protocol designers, for creating protocols that perform close to the capacity, without being overly complicated.

REFERENCES

[1] S. Toumpis, "Capacity bounds for three types of wireless networks: Asymmetric, cluster and hybrid," in *Proc. ACM MobiHoc*, Roppongi, Japan, May 2004.

[2] S. Toumpis, A. J. Goldsmith, and J. Sayir, "Capacity results for asymmetric wireless networks," in *Proc. IEEE International Zurich Seminar*, Zurich, Switzerland, Feb. 2004, pp. 180–183.

[3] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. Inform. Theory*, vol. 46, no. 2, pp. 388–404, Mar. 2000.

[4] R. Negi and A. Rajeswaran, "Capacity of power constrained ad-hoc networks," in *Proc. IEEE INFOCOM*, vol. 1, Hong Kong, China, Mar.-Apr. 2004, pp. 443–453.

[5] C. Peraki and S. D. Servetto, "On the maximum stable throughput problem in random networks with directional antennas," in *Proc. ACM MobiHoc*, Annapolis, MD, June 2003, pp. 76–87.

[6] U. C. Kozat and L. Tassiulas, "Throughput capacity of random ad hoc networks with infrastructure support," in *Proc. ACM MobiCom*, San Diego, CA, Sep. 2003, pp. 55–65.

[7] B. Liu, Z. Liu, and D. Towsley, "On the capacity of hybrid wireless networks," in *Proc. IEEE INFOCOM 2003*.

[8] J. Li, C. Blake, D. S. J. D. Couto, H. I. Lee, and R. Morris, "Capacity of ad hoc wireless networks," in *Proc. ACM MOBICOM*, vol. 1, Rome, Italy, July 2001, pp. 61–69.

[9] E. Perevalov, R. Blum, A. Nigara, and X. Chen, "Route discovery and capacity of ad hoc networks," in *Proc. IEEE GLOBECOM*, St. Louis, MO, Nov.-Dec. 2005.

[10] E. Perevalov, R. Blum, and D. Safi, "On the capacity of ad hoc networks with clustering," in *Proc. IEEE WCNC 2005*.

[11] P. Jacquet and G. Rodolakis, "Multicast scaling properties in massively dense ad hoc networks," in *Proc. International Conference on Parallel and Distributed Systems*, Fukuoka Institute of Technology, Japan, July 2005, pp. 93–97.

[12] M. J. Neely, E. Modiano, and C. E. Rohrs, "Power allocation and routing in multibeam satellites with time-varying channels," *IEEE/ACM Trans. Networking*, vol. 11, no. 1, pp. 138–152, Feb. 2003.

[13] G. Sharma and R. R. Mazumdar, "On achievable delay/capacity trade-offs in mobile ad hoc networks," in *Proc. Workshop on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt)*, Cambridge, UK, Mar. 2004.

[14] S. Toumpis and A. J. Goldsmith, "Large wireless networks under fading, mobility, and delay constraints," in *Proc. IEEE INFOCOM*, Hong Kong, China, Mar.-Apr. 2004.

[15] M. V. Eyoboğlu and J. G. D. Forney, "Trellis precoding: Combined coding, precoding and shaping for intersymbol interference channels," *IEEE Trans. Inform. Theory*, vol. 38, no. 2, pp. 301–314, Mar. 1992.

[16] T. Salonidis, Bhagwat, L. Tassiulas, and R. LaMaire, "Distributed topology construction of Bluetooth personal area networks," in *Proc. IEEE INFOCOM*, vol. 3, Anchorage, AL, Apr. 2001, pp. 1577–1586.

[17] W. Feller, *An Introduction to Probability Theory and its Applications*, 3rd ed. New York, NY: John Wiley and Sons, 1975, vol. I.

[18] R. Motwani and P. Raghavan, *Randomized Algorithms*, 1st ed. Cambridge, UK: Cambridge University Press, 1995.

[19] N. Bansal and Z. Liu, "Capacity, delay and mobility in wireless ad-hoc networks," in *Proc. IEEE INFOCOM*, San Francisco, CA, Apr. 2003, pp. 1553–1563.

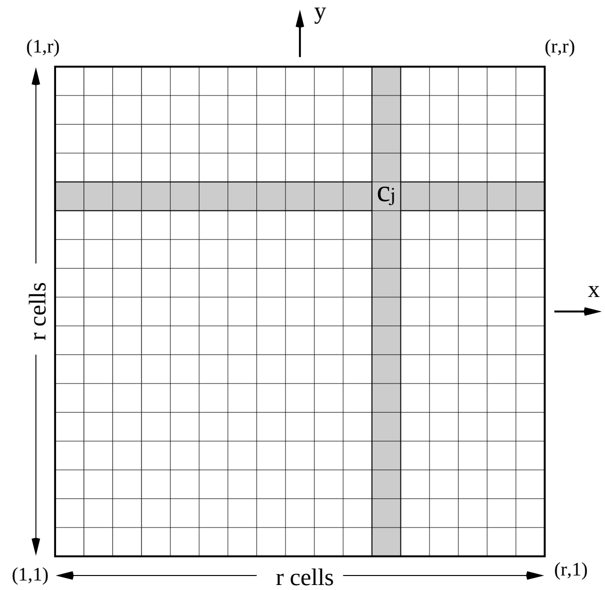

Fig. 1. Partition of the square region $\{(x, y) : |x|, |y| \leq \frac{1}{2}\}$ into a regular lattice of r^2 cells. We define s_j as the number of source nodes in cell c_j, M_j as the number of source nodes lying in cells who share the same x-coordinate with c_j (the shaded cell column) and N_j as the number of destination nodes lying in cells who share the same y-coordinate with c_j (the shaded cell row).

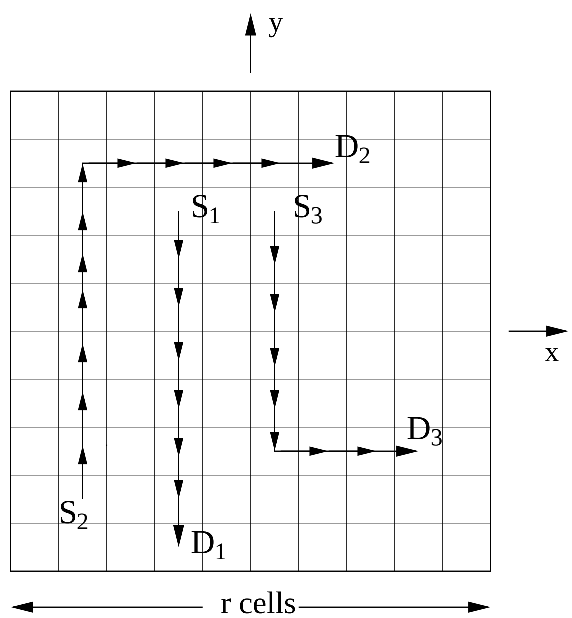

Fig. 2. Examples of routes used in asymmetric networks.

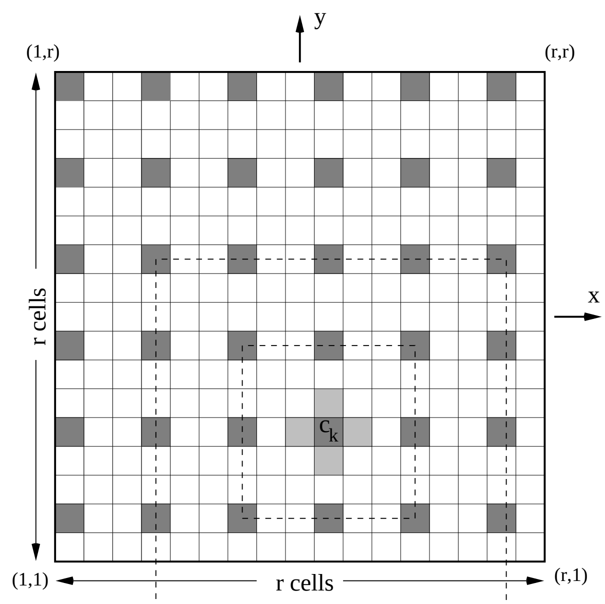

Fig. 3. One of the 9 sub-lattices of cells appears shaded. Only nodes in that sub-lattice are allowed to receive in the corresponding slot, and only from nodes in the same or neighboring cells. The neighbors of cell c_k are lightly shaded. The cells belonging to the same sub-lattice as cell c_k may be placed in at most $\lfloor \frac{r-1}{3} \rfloor$ concentric squares of increasing size, centered at c_k. The first two such squares are denoted by dashed lines.

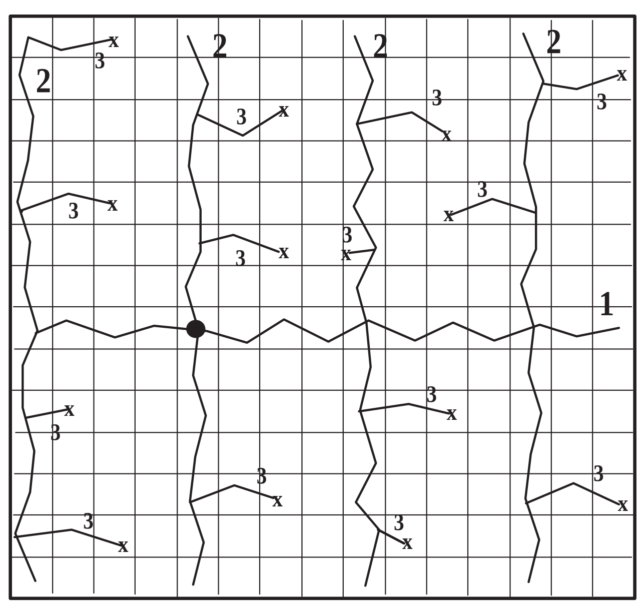

Fig. 4. An example of a multicast tree according to the rules of Section V. The three legs of the tree are denoted by the numbers 1, 2, 3. The source is denoted by a full circle and the destinations by crosses. Relaying nodes are not shown.